

# Bursting the Cloud



Lori MacVittie, 2008-03-09

The cloud computing craze is leading to some interesting new terms. Cloudware and cloudbursting are two terms I particularly like for their ability to describe specific computing models based on cloud computing. Today we're going to look at cloudbursting, which is basically a new twist on an old concept.



Cloudbursting appears to be to marry the traditional safe enterprise computing model with cloud computing; in essence, bursting into the cloud when necessary or using the cloud when additional compute resources are required temporarily. Jeff at [Amazon Web Services Blog](#) talks about the inception of this term as applied to the latter and describes it in his blog post as a method used by Thomas Brox Røst to regenerate a number of dynamic pages in 5 hours rather than the 7 hours that would be required if he had attempted such a feat internally. His approach is further described on [The High Scalability Blog](#).

Cloudbursting can also be used to shoulder the burden of some of an application's processing. For example, basic application functionality could be provided from within the cloud while more critical (e.g. revenue-generating) applications continue to be served from within the controlled enterprise data center. This assumes that only a portion of consumers will actually be interacting with the data-driven side of a web site (customer management, process visibility, etc...) while the greater portion will simply be browsing around on the non-interactive, as it were, side of the site.

Bursting has traditionally been applied to resource allocation and automated provisioning/de-provisioning of resources, historically focused on bandwidth. Today, in the cloud, it is being applied to resources such as servers, application servers, [application delivery systems](#), and other infrastructure required to provide on-demand computing environments that expand and contract as necessary, without manual intervention.

This requires the ability to automate the cloud's data center. Data center automation in a cloud computing environment, regardless of the opacity of the model, requires more than simple workflow systems. It requires on-demand control and management over all devices in the delivery chain, from the storage to the application and web servers to the [load-balancers](#) and [acceleration](#) offerings that deliver the applications to end-users. This is more akin to data center *orchestration* than it is *automation*, as it requires that many moving parts and pieces be coordinated in order to perform a highly complex set of tasks seamlessly and with [as little manual intervention as possible](#). This is one of the [foundational requirements of a cloud computing infrastructure](#): on-demand, automated scalability.

Data center automation is nothing new. Hosting and service providers have long automated their data centers in order to reduce the cost of customer acquisition and management, and to improve efficiency of provisioning and de-provisioning processes.

These benefits can also be realized inside the data center, regardless of the model being employed. The same automation required for smooth, cost-effective management of a cloud computing data center can be utilized to achieve smooth, cost-effective management of an enterprise data center.

The hybrid application deployment model involving cloud computing requires additional intelligence on the part of the [application delivery network](#). The application delivery network must be able to understand what is being requested and where it resides; it must be able to [intelligently route requests](#). This, too, is a fundamental attribute of [cloud computing infrastructure](#): intelligence. When distributing an application across multiple locations, whether local servers or remote data centers or "in the cloud", it becomes necessary for a controlling node to properly route those requests based on application data. In a less sophisticated model, global load balancing could be substituted as a means of directing requests to the appropriate site, a task for which [global load balancers](#) seem a perfect fit.

A hybrid approach like cloudbursting seems to be particularly appealing. Enterprises seem reluctant to move business critical applications into the cloud at this juncture but are likely more willing to assign responsibility to an outsourced provider for less critical application functionality with variable volume requirements, which fits well with an on-demand resource bursting model.

Cloudbursting may be one solution that makes everyone happy.



---

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | [f5.com](http://f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)